

**APPARATUS FOR TRACKING CONNECTION OF SERVICE PROVIDER
CUSTOMERS VIA CUSTOMER USE PATTERNS**

TECHNICAL FIELD

The present invention relates to a method and apparatus for obtaining information on service provider customers and, more particularly, to a method and apparatus for tracking customer connection to a service provider using customer profiles indicating customer use patterns.

BACKGROUND OF THE INVENTION

It is not uncommon for a customer of a service provider, such as an Internet service provider or a long-distance telephone service provider, to obtain and maintain multiple accounts, either simultaneously or at different periods of time. Service providers often find it desirable to match such multiple customer accounts with the single customer. This customer identification enables the service provider to ensure continuity in the type of service provided to the customer, to identify highly valued customers and/or to identify less desirable customers or customers with delinquent accounts. Traditional information, such as name and address, are usually used to perform the customer account matching. This method of account matching suffers from certain disadvantages.

More particularly, it is often difficult for the service providers to perform the account matching on their own existing data. Therefore, service providers usually provide account information to outside vendors who are paid to perform matching against their databases. Since the service providers cannot utilize existing data to perform the matching, they must incur the cost of hiring outside vendors to perform the task. In addition, it is often difficult for the outside vendors to match multiple accounts belonging to a single customer using traditional identifying information, since this information is often entered differently for each account and is subject to frequent errors in data entry. In

sum, this method for customer account matching is costly, and often inaccurate.

Therefore, a method and apparatus for tracking the connection of customers of a service provider are needed which would enable the service
5 provider to easily and accurately track customer movement or connection. The present invention was developed to accomplish these and other objectives.

SUMMARY OF THE INVENTION

In view of the foregoing, it is a principal object of the present invention to
10 provide a method and apparatus that eliminates the deficiencies of the prior art.

It is a further object of the present invention to provide a method and apparatus for accurately tracking movement and/or connection of service provider customers by accurately matching multiple accounts belonging to a
15 single customer.

It is yet another object of the present invention to provide a method and apparatus for accurately matching multiple accounts belonging to a single customer by identifying customers based upon patterns in customer use of the service.

It is a further object of the present invention to provide a method and
20 apparatus for accurately matching multiple accounts belonging to a single customer by comparing the pattern of use of a particular account with the patterns of use for each of the remaining accounts of the service, and identifying multiple accounts as belonging to a single customer when the
25 pattern of use for the particular account substantially matches the pattern of use of at least one of the remaining accounts.

It is a further object of the present invention to provide a method and apparatus for accurately matching multiple accounts belonging to a single customer by comparing the pattern of use of each account in a first sample of
30 accounts being investigated with the pattern of use for each of the accounts

constituting a second sample of accounts of the service provider, and determining that an account in the first sample of accounts and at least one account in the second sample of accounts belong to a single customer when the pattern of use for the account in the first sample of accounts substantially
5 matches the pattern of use of at least one of the accounts in the second sample of customers.

It is a further object of the present invention to provide a method and apparatus for accurately matching multiple accounts belonging to a single customer by comparing the pattern of use of each of the accounts in a first
10 sample of accounts being investigated with the patterns of use for each of the accounts constituting a second sample of accounts of the service provider, where the second sample of accounts constitutes a subset of all of the accounts of the service provider, and determining that an account in the first sample of accounts and at least one account in the second sample of accounts belong to
15 a single customer when the pattern of use of the account substantially matches the pattern of use of at least one of the accounts in the second sample of customers.

It is yet another object of the present invention to provide a method and apparatus for assigning customer history information to multiple accounts
20 belonging to a single customer by comparing the pattern of use of a particular account of the single customer with the pattern of use for the remaining accounts, identifying multiple accounts as belonging to the single customer when the pattern of use for the particular account substantially matches the pattern of use of at least one of the remaining accounts, and assigning the
25 customer history information of the matching remaining account(s) to the particular account.

These and other objects and features of the present invention will be apparent upon consideration of the following detailed description of preferred embodiments thereof, presented in connection with the following drawings in
30 which like reference numerals identify like elements throughout.

BRIEF DESCRIPTION OF THE DRAWINGS

In the drawings,

Fig. 1 illustrates an example of a telecommunications system;

5 Fig. 2 illustrates the basic elements of a system for performing the method according to the present invention;

Fig. 3 illustrates a flow diagram of the steps performed in the method according to the present invention;

10 Fig. 4 illustrates an example of customers of a service provider and the information obtained from the customers for performing the method according to the present invention;

Fig. 5 illustrates an example of how one customer account is distinguished from the remaining customer accounts in a sample according to the present invention; and

15 Fig. 6 illustrates another example of how of how one customer account is distinguished from the remaining customer accounts in a sample according to the present invention.

DETAILED DESCRIPTION OF THE INVENTION

20 In order to facilitate the description of the present invention, the invention will be described with respect to the particular example of long-distance telephone service providers. Examples will be described that illustrate particular applications of the invention for long-distance telephone service. The present invention, however, is not limited to any particular service provider
25 nor limited by the examples described herein. Therefore, the description of the embodiment that follows is for purposes of illustration and not limitation.

A particular application of the present invention is identifying customers of a long-distance telephone service provider who have moved without informing the long-distance telephone service provider. The identification is
30 accomplished by matching the customer's pre-move and post-move calling

patterns. For example, upon moving, a customer usually notifies their local telephone service provider to disconnect their service and close their account, and reconnect and open another account at the new location. The customer often assumes that their long-distance service provider will be notified and will be able to match their new account with their past history. However, the information provided to the long-distance provider from the local service provider does not distinguish a "mover" from a "new connect". Therefore, the long-distance provider cannot match the new account with the past history of the customer.

Another application of the invention with respect to long-distance telephone service providers is to identify a customer who has opened multiple accounts with different names or where the names have been entered differently by a data entry person (e.g., one account under Joseph Smith and another account under Ellen Smith or one account under Joseph Smith and another account under J. Smith). The present invention solves these and other problems by using customer calling pattern data to match multiple accounts belonging to a single customer.

Referring to Fig. 1, an exemplary telecommunications network 10 is shown. Local switching offices 12 and 14 are connected to each other by trunk 20, while local switching offices 16 and 18 are connected to each other by trunk 22. Trunk 20 is used to route calls from a telephone 26 served by the local switching office 12 to a telephone 28 serviced by the terminating local switching office 14. Long-distance calls to telephone 32, for example, are processed by a long-distance network 30. Service within local access and transport areas for local calls is often provided by local telephone exchange carriers such as Bell South, and service between the local access and transport areas for long-distance calls is often provided by interexchange carriers such as AT&T.

Fig. 2 illustrates the basic elements of a system for performing the method according to the present invention. A data processing system 34

receives the customer use detail information for each customer via interface 33. The data processing system 34 uses the use detail information for each customer to generate customer profiles to distinguish customers from one another. The data processing system 34 outputs a profile for each service provider customer based upon service use patterns. In addition, it performs the account comparison step and outputs probable matches based upon the result of the comparison, as well as the overlapping information associated with the probable match. The output from the data processing system 34 may be supplied to a monitor 35, printer 36, and/or other output device 37. The customer use data information may be obtained by any appropriate means utilized by the service provider.

According to the example of the present invention described herein, customers are characterized by the calls they make on the long-distance network 30. The long-distance telephone service provider records call information for billing purposes. This data is recorded at a very low error rate. This data may include the telephone number making the call, the telephone number being called, the time of day, and/or the duration of the call, for example. This existing data may be used as call detail information to determine the call patterns of the service provider customers. The call patterns may then be used to identify multiple accounts belonging to a single customer. Therefore, no outside vendor services are required to perform the present invention.

A customer's calling use is, of course, variable, but the present invention exploits the fact that people tend to call certain numbers (e.g., family) repeatedly over time. By identifying customers or accounts by groups of called numbers, the method according to the present invention is robust to variations in calling pattern over time. The general procedure is described below with reference to Fig. 3.

First, depending on the purpose of the matching, a sample of known customers of interest is identified in steps S1 and S2. The sample may include

all of the service provider customers or a subset of the service provider customers. For example, if the service provider is trying to detect when a new account belongs to a high-value customer who recently moved without informing the service provider, the sample consists of high-value customers
5 who disconnected from the service provider at about the same time that a new account(s) was opened. Then, the service provider obtains the list of all calls to or from the customers in the sample over a chosen time period, e.g., one or two months, as shown in step S3. In step S4, the service provider aggregates any customer calls to or from a single number into a "feature." The service provider
10 eliminates from the set of features all calls that are relatively frequent across the sample with respect to an experimentally chosen threshold. For example, if many of the customers in the sample called a major catalog merchant, the merchant's number has little power for distinguishing one customer from another, and it is ignored. The next step is to construct a profile or
15 "fingerprint" for each customer in the sample, i.e., a subset of the features attributed to the customer, so that any two customers in the sample have different profiles (although their profiles may have some features in common), as shown in step S5.

Feature selection and profiling are also performed for each of the
20 accounts being investigated (continuing with the movers example, the set of new accounts connected to the service provider), as shown in steps S6-S10 for identifying customers in a second sample of accounts to be later compared to the first sample in steps S1-S5. The comparison for account matching of the first sample of accounts to the second sample of accounts is performed in step
25 S11 to look for accounts in the two samples that have similar fingerprints. In step S12, probable matches are determined in accordance with the results of the comparison in step S11.

The method according to the present invention is not hindered by minor changes in customer use, because the method does not require identical
30 fingerprints for a match. The matched accounts may not be for the same

customer, but they have a relatively high probability of being so, and are cost-effective targets for further investigation, such as by contacting the customer directly.

The method and apparatus according to the present invention provides
5 numerous advantages to the service provider. More specifically, the information obtained by the matching method and apparatus may be used to provide continuity of service to customers and to identify high-value customers. In addition, the information may be used to screen new accounts for creditworthiness. If a new customer's call pattern matches an old account with
10 a history of delinquency or fraud, the service provider would have the opportunity to restrict the new account, avoiding the generation of new uncollectible debt. The service provider may also use the information to pursue collection of an old account's debt, gaining revenue that would otherwise be lost.

15 A more detailed description of the method and apparatus according to the present invention as implemented in the telecommunications example is set forth below.

PHASE I. PROFILING CUSTOMERS

20 A long-distance customer's calling activity includes a sequence of inbound and outbound calls. Those calls have a number of attributes, such as duration, time of day, time of week, etc. If a customer's calling activity is observed over some time window, additional attributes can be derived. Derived attributes may include frequency of calls made by the customer to a particular
25 number, average talk time, total talk time, etc. It will be appreciated by those of ordinary skill in the art that many different attributes and derived attributes are possible.

In the example illustrated in Fig. 4, customers A and B are subscribers of a particular long-distance telephone service provider. Customer C is not a
30 subscriber to the particular long-distance provider. Assume that customer A is

a member of a first sample of customers representing the customers to be investigated. The first sample may be a subset of all of the service provider customers. For example, assume that the first sample of customers includes high value customers who disconnected from the service provider around the time that a new connection(s) occurred. Also assume that the profiles of the customers in this sample, including the profile of customer A, will be matched or compared against the profiles of the service provider customers constituting a second sample of customers. In the present example, assume that the second sample represents new account(s) of the service provider.

In order to create a profile for customer A, customer A's call pattern for some period of time is monitored. The period of time may be selected by the service provider and need only be long enough to observe a calling pattern. Step S20 in Fig. 4 illustrates the step of recording information for each of customer A's calls. In the current example, the information recorded for customer A includes the number making the call, the number being called, the time of day, and the duration of the call. Of course, the recorded information may include more or less information or different information than that shown in the example of Fig. 4.

A profile for customer A is generated from the recorded information, as shown in step S21. More particularly, for each of customer A's calls, the call's frequency within the period of time may be determined. In addition, each call may be distinguished as either an inbound call or an outbound call. In this framework, customer A's call detail may be defined by a set of features (phone numbers), each determined uniquely by a number called and a direction (inbound/outbound). More formally, a feature of customer A may be defined to be the pair:

$$f = (c, I/O \text{ indicator})$$

where f = a feature (Billed Telephone Number (BTN)); c = number (BTN) called by A (outbound call) or the number (BTN) which called A (inbound call); and the I/O indicator = 1 for inbound calls or 0 for outbound calls. Also, let the frequency ϕ_{Af} be the ratio of calls A made (received) to (from) c relative to all calls A made (received).

Customer A is identified by its BTN. Calls c are identified by their BTNs as well. If an outbound call is made by A to c , c doesn't even have to be a customer of the service provider. For example, in Fig. 4, customer A may place a call to customer C, who is not a customer of the particular service provider (i.e., customers A and C subscribe to different long-distance telephone service providers). All calls passing through A's service provider's network are observed. Therefore, the data for creating the customer profile will include all calls made by customer A, even if the recipient is not a customer of the service provider. The reverse is not true, however, because a call from a non-subscriber may reach customer A without touching the service provider network to which customer A subscribes. Therefore, for inbound calls, only calls coming from the customers of the service provider to which A subscribes are observed. In Fig. 4, this means that only calls from customer B to customer A are observed, as shown in Step S21. The steps shown in Fig. 4 correspond to steps S1-S4 in Fig. 3.

The list of features of customer A can be large, consisting of all calls that "touched" customer A in the given time window. Some of those calls may be incidental, while some may be part of a consistent pattern specific to customer A. The goal of profiling is to substantially reduce the number of features to be included in customer A's profile, and still be able to differentiate between customer A and all other customers in the sample. Therefore, features that are incidental may be eliminated. This may be accomplished by selecting a frequency threshold τ_f , and including feature f in the list of features for customer A only if the frequency ϕ_{Af} exceeds the frequency threshold τ_f . In calculating the threshold τ_f , only calls to/from customers in the sample of

profiled customers are considered. Therefore τ_f is based on the sample data. The threshold can be selected based upon experience of the service provider, on experiments performed by the service provider or based on training. The threshold varies depending upon the service provider and on the needs of the service provider. Therefore, any suitable method for selecting the threshold may be used. The goal is to create a calling *profile* or "*fingerprint*" for customer A, as shown in step S5 in Fig. 3, such that:

- (i) the profile allows the service provider to distinguish customer A from all the other customers in the sample;
 - (ii) the profile consists of as few features as possible;
 - (iii) the features in the profile either are unique to customer A or appear on the feature list of as few other customers as possible; and
 - (iv) the features in the profile are with high probability a part of customer A's typical pattern.
- Step (iv) can be accomplished by limiting feature inclusion with appropriate thresholds on ϕ_{Af} , while steps (i), (ii) and (iii) can be accomplished with an integer optimization model of set covering, as discussed below.

Set Covering Problem According to a First Embodiment

- For any set S, let $|S|$ represent the number of members of S. Let the set Ω be the first sample of customers, with $|\Omega| = m + 1$. Let $A \in \Omega$ be a customer to be profiled, and let F_A be the list of all features f of customer A whose frequencies ϕ_{Af} pass their associated thresholds τ_f . Let $n = |F_A|$. Consider customer $B \in \Omega$, $B \neq A$. Customer A's pattern is determined to differ significantly from customer B's with respect to feature f if one of the two passes threshold τ_f and the other does not, i.e.,

$$A \neq B \text{ if and only if } \phi_{Af} > \tau_f \text{ and } \phi_{Bf} \leq \tau_f \text{ (or vice versa).}$$

For every feature f from the list F_A , create a set $S_f = \{B: \phi_{Bf} \leq \tau_f\}$. S_f is the subset of Ω such that $B \in S_f$ if and only if $A \neq_f B$. Then $\bigcup_f S_f$ is the set of all customers in Ω that differ significantly from customer A on at least one feature. Assume that $\bigcup_f S_f = \Omega \setminus A$; that is, no customer other than customer A has a list

5 of features identical with customer A . It follows, that $|\bigcup_f S_f| = m$. The problem of minimizing the number of features used to distinguish customer A can be formulated as minimizing the number of sets S_f needed to cover $\Omega \setminus A$, as follows:

10 Let $a_{Bf} = 1$ if $B \in S_f$, and 0 otherwise.
Let $x_f = 1$ if f is selected, and 0 otherwise.

$$\text{minimize } \sum_{f=1}^n x_f$$

$$\text{subject to } \sum_{f=1}^n a_{Bf} x_f \geq 1 \quad \text{for } B = 1, \dots, m$$

15 and $x_f = 0$ or 1 for $f \in F_A$.

Given the optimal solution x^* , cover $F^* = \{f: x_f^* = 1\}$ provides a minimal set of features establishing customer A as distinctly different from each customer B in the sample to be investigated. This integer program is an example of the set
20 covering problem.

Greedy Algorithm

The set covering problem can be solved for each customer. When there are a large number of customers (e.g., 1.5 million) disconnecting each month,
25 computational efficiency is essential, so a global optimum may not be sought.

The set covering problem can be efficiently solved by a greedy heuristic as follows:

- Step 0: $F^* = \emptyset; \Omega^* = \Omega \setminus A.$
- 5 Step 1: Select f such that $|S_f| = \max\{|S_j| : j \in F_A \setminus F^*\}$
- Step 2: $F^* = F^* \cup \{f\}; S_j = S_j \setminus S_f$ for all $j \in F_A \setminus F^*; \Omega^* = \Omega^* \setminus S_f;$
- Step 3: If $\Omega^* = \emptyset$, STOP and output cover F^* . Otherwise, go to Step 1.

Generalized Set Covering Problem According to the Second Embodiment

- 10 According to this embodiment, a generalized set covering problem is used to accomplish steps (i), (ii) and (iii), noted above with respect to profiling or fingerprinting. The *generalized set covering* formulation looks for a cover in which every element is covered k times (belongs to at least k sets). k may be defined as the *depth* of the cover. This model may be more appropriate for
- 15 obtaining a robust profile, so that a customer whose call pattern changes to some degree following a move can still be recognized. Multi-covering allows selection of more features than a regular set covering. In addition, different customers can be covered to different depths when appropriate. For example, it may be desirable to cover customer B_1 , who has a large number of features,
- 20 more times than customer B_2 , whose list of features is much smaller.

For each customer B , let k_B be the desired depth of coverage. Then the multicover formulation is as follows:

$$\begin{aligned}
 & \text{minimize} && \sum_{f=1}^n x_f \\
 25 & \text{subject to} && \sum_{f=1}^n a_{Bf} x_f \geq k_B \quad \text{for } B = 1, \dots, m \\
 & \text{and} && x_f = 0 \text{ or } 1 \quad \text{for } f \in F_A
 \end{aligned}$$

Cover $F^* = \{f: x_f = 1\}$ provides a minimal set of features establishing customer A as distinct from each customer B in the sample to be profiled at the level of coverage desired for robustness.

5 Modified Greedy Algorithm:

Step 0: $F^* = \emptyset; \Omega^* = \Omega \setminus A$

Step 1: Select f such that $|S_f| = \max \{|S_j|: j \in F_A \setminus F^*\}$

Step 2: $F^* = F^* \cup \{f\}$

Step 3: For all $B \in S_f$:

10 $k_B = k_B - 1;$

if $k_B = 0$, $\Omega^* = \Omega^* \setminus B$ and $S_j = S_j \setminus B$ for all $j \in F_A \setminus F^*$

Step 4: If $\Omega^* = \emptyset$, STOP and output cover F^*

If $F^* = F_A$, STOP; no cover of depth k exists.

Otherwise, go to Step 1.

15

In Fig. 5, a feature f is shown with a line connecting it to a member of the sample (A, B1, B2, B3) only if f passes the threshold for that member. Customer A, the customer being profiled, is connected to all of the features. Thus, customer B1 can be covered, for example, once by choosing any feature that has no connecting line to B1. In the fingerprint of customer A pictured in
20 Fig. 5, coverage for each of the customers B1, B2, B3 is at depth 3. The greedy algorithm may first include in the profile the two features (5 and 7) that are unique to customer A. This gives each of customers B1, B2, and B3 coverage of depth 2. Next, feature 1 may be included in the profile. It covers customers
25 B2 and B3, which now have coverage of depth 3. Finally, feature 2 is included, which covers customers B1 and B3. Now the depth of coverage for customers B1 and B2 equals 3, while the depth of coverage for customer B3 equals 4. In this example, customer A has a fingerprint of size 4 (marked by thicker lines).

If a customer being fingerprinted has at least K unique features (that is,
30 features which do not appear in the call detail of the remaining customers in

the sample being profiled) and $K \geq \max_B \{k_B\}$, then the set covering model may not be necessary, because a collection of K unique features can be used for a profile. In the example pictured in Fig. 6, the fingerprint of customer A consists of 5 features, all of which are unique to customer A. Customers B1, B2, B3 are
5 all covered at the depth of 5. This saves computing time by applying the set covering algorithm only in cases where there are not enough unique features.

The customer profile should be stable to be of use. In the present example, the customer profile may be considered as stable if, for example, a fingerprint including at least 5 features can be found in each month, and either
10 at least 30% of the features remain unchanged from month to month or at least three features remain unchanged from month to month. However, the stability of a fingerprint can be defined as appropriate for the particular service provider.

15 PHASE II: MATCHING CUSTOMERS TO PROFILES

Referring again to Fig. 3, in steps S1 and S2, the phone numbers of customers to be profiled, such as customer A in Fig. 4, are determined. In step S3, call details of calls over the service provider network involving these customers are determined over a predetermined period of time. In step S4,
20 calls for a particular number appearing more than once are aggregated into a feature, and in step S5 the appropriate algorithms, as discussed above, are applied to choose features that uniquely identify as many customers as possible of the customers to be profiled. In steps S6 and S7, the phone numbers of customers to be recognized or new connect customers are
25 determined. In step S8, the call details of calls over the service provider network involving new connect customers are determined over a predetermined period of time. This predetermined period of may be the same, less than or more than the time period in step S3. In step S9, calls for one number are aggregated into a feature, and in step S10, the appropriate algorithms, as
30 discussed above, are applied to choose features that uniquely identify as many

customers in this group as possible. In step S11, the profiles of the customers from the second sample of customers being investigated are compared with the profiles of the first sample of customers, including customers A and B. When a profile in one group overlaps a profile in the other group, it is determined that
5 there is a strong possibility of a match in step S12, indicating that the profiles correspond to accounts belonging to the same customer. The service provider may act on results indicating a strong possibility of a match by further investigating the particular accounts and possibly contacting the customer directly.

10 According to the present invention, as described above with respect to the example of a telecommunications system and a long-distance telephone service provider, multiple accounts belonging to a single customer may be easily and accurately determined by the service provider. Therefore, according to the present invention, accurate results can be obtained without requiring
15 the services of an outside vendor.

While particular embodiments of the invention have been shown and described, it is recognized that various modifications thereof will occur to those skilled in the art without departing from the spirit and scope of the invention.